HP OpenView

Storage Mirroring application notes

Guidelines for networking and failover

Part number: T2558-88080
Second edition: June 2006

Legal and notice information

Storage Mirroring Guidelines for networking and failover application notes

Document overview

This document is a Storage Mirroring application note. An application note provides guidelines on the use of Storage Mirroring in a specific environment.

This document contains:

- **Document overview**—Explains what an application note contains, how it should be used, what you need to know before trying to use the application note, and where you can go for more information.
- **Implementation guidelines**—Describes all of the considerations that you must weigh when implementing your Storage Mirroring solution. Includes both basics, such as system requirements, as well as configuration and environment-specific topics, such as interactions with specific clients or special considerations for WAN (Wide Area Network) environments. Pay special attention to those topics that are directly related to your environment.

Audience

This document is written for network and application administrators who have a working understanding of the applications and environments where the Storage Mirroring solution is to be deployed. You many need to expand on the documented information in order to customize the solution to fit your environment.

Before you use this application note, you should have an understanding of basic Storage Mirroring operation.

Expectations

Application notes are intended to provide a framework for configuring a Storage Mirroring solution in a specific environment and to draw attention to decisions you will need to make when configuring your solution.

Because there are an infinite number of possible configuration, network, and environment scenarios, application notes contain general configuration guidelines as well as an example configuration procedure that has been tested for a specific environment.

This document assumes that you are comfortable working with your operating system, Storage Mirroring, and your application(s).

Related documentation

Before you begin to configure your solution, make sure that you have complete documentation for your operating system, application, and Storage Mirroring. This application note does not provide step-by-step instructions for using standard operating system, application, and Storage Mirroring functionality.

The following documents contain additional information that you may need while setting up this solution:

- *HP OpenView Storage Mirroring user's guide* or online documentation

Getting help

For help using Storage Mirroring, refer to the Storage Mirroring online manual or online help.

Implementation Guidelines

It is helpful to understand how Windows systems communicate when implementing a high availability solution using Storage Mirroring failover capabilities. Storage Mirroring failover can be implemented with very little configuration necessary in small or simple networks, but additional configuration may be required in large or complex environments to ensure that all clients will resolve the source name to the correct IP address.

Understanding how Server Message Block (SMB) communications occur in Windows NT 4.0, Windows 2000, and mixed environments is the key to understanding what actions must be taken at failover in regard to Domain Name System (DNS) and Windows Internet Naming Service (WINS) services. Due to the way SMB is implemented in Windows 2000, pure Windows 2000 environments should require no name resolution adjustments after failover. However, in Windows NT 4.0 or mixed environments, successful SMB communications require specific name resolution requirements to be met.

This document also discusses other networking topics related to failover, including how to view name caches and the Address Resolution Protocol (ARP) cache to troubleshoot. Additionally, common questions such as how failover affects domain controllers and how to fail over Internet Protocol (IP) addresses to remote targets are addressed.

Storage Mirroring failover

When Storage Mirroring failover occurs, the following events take place. (Depending on the failover options selected, the events listed may not occur or may occur in a different order.)

- The pre-failover script runs.
- The specified source IP addresses are added to the target network adapters.
- An unsolicited ARP is sent to inform the network segment of the change in IP address-to-Media Access Control (MAC) address mappings.
- The source's DNS host name is added to the system as a NetBIOS name. First Storage Mirroring checks the hosts file and uses the first name there. If there is no hosts file, Storage Mirroring uses the first name in DNS. (Although, the first name in DNS may not always be the same each time the DNS server is rebooted.) Lastly, if there is no DNS server, Storage Mirroring uses the Failover Control Center monitor name.

NOTE: If the DNS host name you want added to the system is not being used, you will need to reconfigure DNS so that only one host entry is available for the monitored IP address. If other names need to resolve to that server, add alias records for those entries.

- The Microsoft Networking Client updates the WINS server, associating the source's host name as returned by DNS with the target's primary IP addresses.
- The source's host Service Principal Names (SPNs) are removed from the source computer account and added to the target computer account.
- The post-failover script runs.

Any other actions beyond the scope of these events that are required for successful failover must be scripted and executed in the pre-failover or post-failover scripts. Ensuring that all clients will be able to access the source after failover often requires failover scripting to update the name resolution servers.

Storage Mirroring failback

When Storage Mirroring failback occurs, the following events take place. (Depending on the failback options selected, the events listed may not occur or may occur in a different order.)

- The pre-failback script runs.
- The source’s IP addresses are removed from the target’s network adapters.
- An unsolicited ARP is sent to inform the network segment of the change in IP address-to-MAC address mappings.
- The source’s DNS name is removed.
- The Microsoft Networking Client updates the WINS server by tombstoning the source’s entry.
- The source host SPNs are removed from the target computer account and added to the source computer account.
- The post-failback script runs.
- The source post-failback script runs.

Any other actions beyond the scope of these events that are required for successful failback must be scripted and executed in the pre-failback or post-failback scripts.

SMB and NetBT

Windows file and print sharing uses the SMB protocol, which has historically relied on NetBIOS. NetBIOS, in turn, required NetBIOS over TCP/IP (NetBT) to function on IP networks. NetBT uses Transmission Control Protocol (TCP) port 139 and has a limitation of binding only to the primary IP address of each Network Interface Card (NIC). This is explained in Microsoft Knowledge Base article 131641 and can be seen by

using a port scanner to probe TCP port 139 (the "nbsession" port) on an adapter with multiple addresses. This will show that NetBT is listening on TCP port 139 only on the primary address.

Windows 2000 and later versions do not require the NetBT layer and use SMB directly on top of TCP/IP using port 445 (TCP and User Datagram Protocol (UDP)). This implementation does not have the aforementioned binding limitation and allows clients to establish SMB sessions to any IP address on the server using port 445. In order to be backward compatible with legacy clients and servers, Windows 2000 also supports SMB on NetBT using port 139, which inherits the primary IP address limitation. If NetBT is disabled, a Windows 2000 system will use only port 445 for SMB session. See the following three diagrams.

Windows NT 4.0 and Earlier

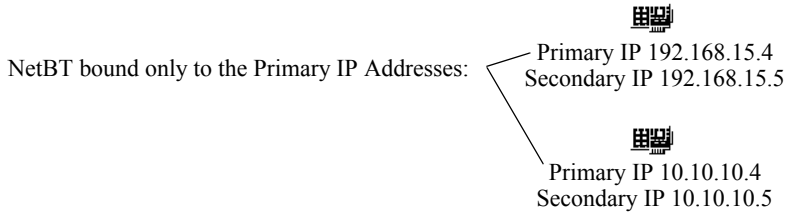

Windows 2000 and Later with NetBT Enabled

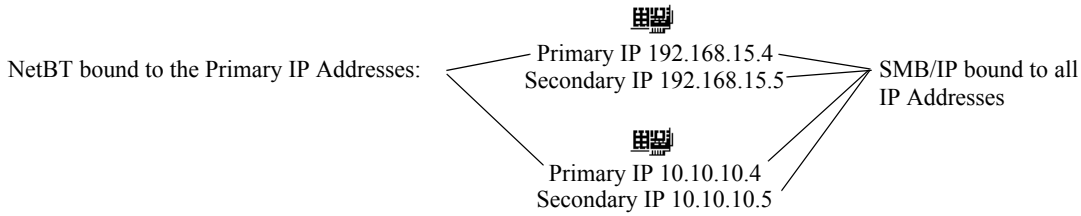

Windows 2000 and Later with NetBT Disabled

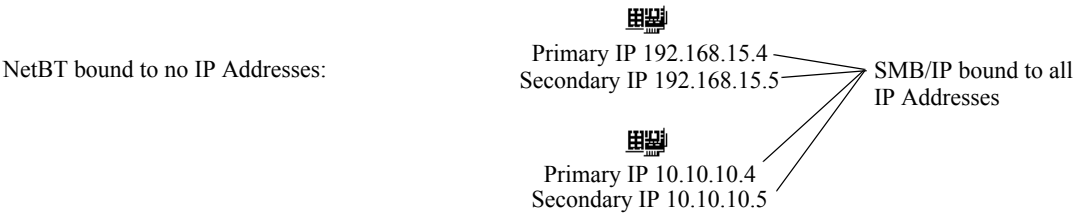

The command `netstat -a -n` will show a Windows 2000 system listening on UDP 0.0.0.0:445 and TCP 0.0.0.0:445. An IP address of 0.0.0.0 indicates a global binding, so it is listening on TCP and UDP port 445 on all IP addresses for SMB communications. The NetBT binding will be shown as TCP 172.16.137.5:139 for the primary IP address on each NIC, indicating that it is listening only on those IP addresses.

Within Microsoft there is not a standardized way to refer to the implementation of SMB on port 445. Knowledge Base article 150543 refers to it as "Direct Hosting of SMB Over TCP/IP." Article 204279 refers to it as the "NetBIOS-less" transport, even while mentioning that the NET CONFIG SERVER command will report the binding as `NetbiosSmb (000000000000)`. Regardless, this document will refer to it as SMB/IP for the sake of convenience.

The NetBT bindings are displayed with the network adapter's Globally Unique Identifier (GUID) and MAC address as:

```
NetBT_Tcpip_{030319AA-84D2-481D-8E28-1EAC6D4AF2A8} (000102d03b7d)
```

The `netstat -a` command returns display names for ports 139 and 445 as `netbios-ssn` and `microsoft-ds` respectively.

Due to the differences between operating system versions, behavior of client/server communications after Storage Mirroring failover will vary depending on the operating system version of the clients and server. When Storage Mirroring failover occurs, the source IP address is added as a secondary IP address by default. Accordingly, if either the target server or clients are pre-Windows 2000 systems, one of the following is required in order for clients to access files on the target using the source's NetBIOS name:

- WINS and DNS servers must be updated so the source NetBIOS and host name resolve to a primary IP address on the target
- The IP address of the source must be made a primary IP address on the target (the "replace" option)

If all clients and the target server are Windows 2000 or later, WINS and DNS do not need to be updated if the source IP address is failed over. The clients, using SMB/IP, can connect to the target server with either the source or target IP address, so it does not matter which IP address the source name resolves to.

Microsoft Knowledge Base article 142309 lists the order of name resolution methods used by clients when resolving a NetBIOS name:

1. NetBIOS name cache
2. WINS server
3. B-node broadcast
4. LMHOSTS file
5. HOSTS file
6. DNS server

Failover and WINS

When Storage Mirroring failover occurs, Windows becomes aware of the new name that the server now owns and initiates a WINS registration with the target's primary WINS server. This registration associates the source name with the target's primary IP addresses and will be distributed to other WINS servers on the network via WINS replication. The length of time required for all WINS servers to obtain the new registration will depend on the number of WINS servers, the architecture of WINS replication, and the interval of replication. If there is only one WINS server or if the target and all clients that need access are configured with the same primary WINS server, then no additional configuration is necessary.

However, if some or all clients are configured with primary WINS servers other than the target's primary WINS server, failover scripts can be used to make the necessary WINS registrations on all WINS servers or to initiate WINS replication. The first method, making the WINS registrations, incurs less network overhead, but will require the appropriate permissions (administrator group membership) on all WINS servers. The second method, initiating WINS replication, will rely on the WINS infrastructure to distribute the new records, but will require the system and network resources required to complete WINS replication. The impact of WINS replication will depend on the size of the network and the WINS architecture.

WINS management scripting

WINS registrations can be made via scripts as part of the failover process. This can be accomplished by using the Windows Resource Kit WINS Administration Tool (`winscl.exe`) or the Windows 2000 NETSH command in the failover script to either initiate WINS replication or make the appropriate registration with each WINS server.

WINS scripting is required when some or all clients have a primary WINS server other than the target's primary WINS server, the target server or any clients are pre-Windows 2000 systems, and the "replace" failover option is not used.

WINS scripting is not required when the target and all clients have the same primary WINS server (regardless of whether clients and target are in a LAN or WAN environment), the target server and all

clients are Windows 2000 or later versions and the IP address is failed over, or the “replace” failover option is used.

Since Windows 2000 clients can use SMB/IP with any IP address on a Windows 2000 server, it does not matter if they resolve the source name to the source IP address or the target IP address. Both will work as long as the source IP address is failed over to the target. However, if the source IP address is not failed over (typically because the source and target are on different subnets), WINS servers must be updated at failover so that clients will resolve the source name to the target IP address.

Management of Windows NT 4.0 WINS servers can only be scripted with WINSCL, which includes two different methods of registration. The first way is to import an LMHOSTS file that is located on the WINS server. This method requires less scripting, but more files to manage since a separate LMHOSTS file must be maintained on each WINS server for each source server. If there are ten source servers and ten WINS servers, 100 LMHOSTS files, ten on each WINS server, will be required.

The second method is to script each registration individually. Although this requires only one script per source and does not require any files to be stored on the WINS servers, the scripts are much longer since a complete WINS registration for a member server actually has three registrations: SERVER[0x0], SERVER[0x3], and SERVER[0x20]. See Microsoft Knowledge Base article 119495 for descriptions of names registered with the WINS service. (The SERVER[0x3] entry is a registration for the Messenger service and can typically be omitted unless there is an application dependency.) See ”Creating WINSCL scripts” on page 7 for sample WINSCL scripts.

Windows 2000 WINS servers can be managed with the WINSCL utility or the Windows 2000 NETSH command, which supports numerous Windows 2000 IP management functions. NETSH can be used interactively or in scripts. The “add name” command in the WINS context will register SERVER[0x0], SERVER[0x3], and SERVER[0x20] with the specified WINS server. See Microsoft Knowledge Base article 233375 for more information on adding WINS registrations with the NETSH command. Following is an example that adds a dynamic (RecType=1) registration for a server named EUROPA with an IP address of 172.16.137.5 to a WINS server with an IP address of 172.16.137.1.

```
netsh wins server 172.16.137.1 add name Name=EUROPA RecType=1 IP={172.16.137.5}
```

Using the NETSH command is the preferable method in Windows 2000 environments since each registration is simply a one-line command. Additionally, the NETSH command can be used to initiate WINS replication on Windows 2000 WINS servers if that is the method selected to propagate the WINS changes after failover.

```
netsh wins server [ip address of WINS server] set replicateflag 1
```

Creating WINSCL scripts

Updating WINS servers with WINSCL scripts will require two LMHOSTS files on each WINS server that need to be updated, as well as text files containing the WINSCL commands that will be executed at failover and failback. An LMHOSTS file (TARGET_HOST) mapping the source's name to the target's IP address will be imported into the WINS database at failover, and another file (SOURCE_HOST) mapping the source's name to the source's IP address will be imported into the WINS database at failback.

The following sample LMHOSTS files use a source computer named PRODSVR with an IP address of 10.5.0.2, while the target server's IP address is 10.4.0.4. As you can see, the SOURCE_HOST file associates PRODSVR with its own 10.5.0.2 IP address, while the TARGET_HOST associates PRODSVR with the target's IP address of 10.4.0.4

NOTE: The sample scripts provided are only examples. Because no two environments or configurations are exactly the same, you MUST modify the sample scripts in order to make the solution work in your environment.

SAMPLE_SOURCE_HOST

10.5.0.2	PRODSVR

SAMPLE_TARGET_HOST

```
10.4.0.4                                PRODSVR
```

There are seven WINSCL commands that must be scripted to import an LMHOSTS file into a WINS database. The following example, TARGET.DAT, includes the commands necessary to connect to a WINS server with an IP address of 10.5.0.11 and import the TARGET_HOST file that is located on the WINS server.

SAMPLE_TARGET.DAT

```
1
10.5.0.11
SI
1
D:\SOURCE_HOST
0
EX
```

For this example, the following command would be placed in the post-failover script to run the TARGET.DAT script.

SAMPLE_POST_OVER.BAT

```
WINSCL < D:\SCRIPTS\TARGET.DAT
```

The SOURCE.DAT script would be used at failback to import the SOURCE_HOST file.

SAMPLE_SOURCE.DAT

```
1
10.5.0.11
SI
1
D:\SOURCE_HOST
0
EX
```

The post-failback script would contain the following command to execute the script:

SAMPLE_POST_BACK.BAT

```
WINSCL < D:\SCRIPTS\SOURCE.DAT
```

NOTE: The LMHOSTS files must be located on the WINS server. The WINSCL utility does not actually import the LMHOSTS. It simply tells the WINS service to import the file. Accordingly, the WINS service interprets the path to the file (D:\SOURCE_HOST, for example), and the file must be at that location on the WINS server.

See the Failover chapter of the HP OpenView Storage Mirroring user's guide for information on how to configure failover scripts and the Windows NT 2000/4.0 Resource Kit documentation for WINSCL documentation.

Failover and DNS

DNS resolution is also a consideration after failover, especially when the source IP address is not failed over to the target. Use the following decision tree to determine when DNS host entries need to be changed after failover.

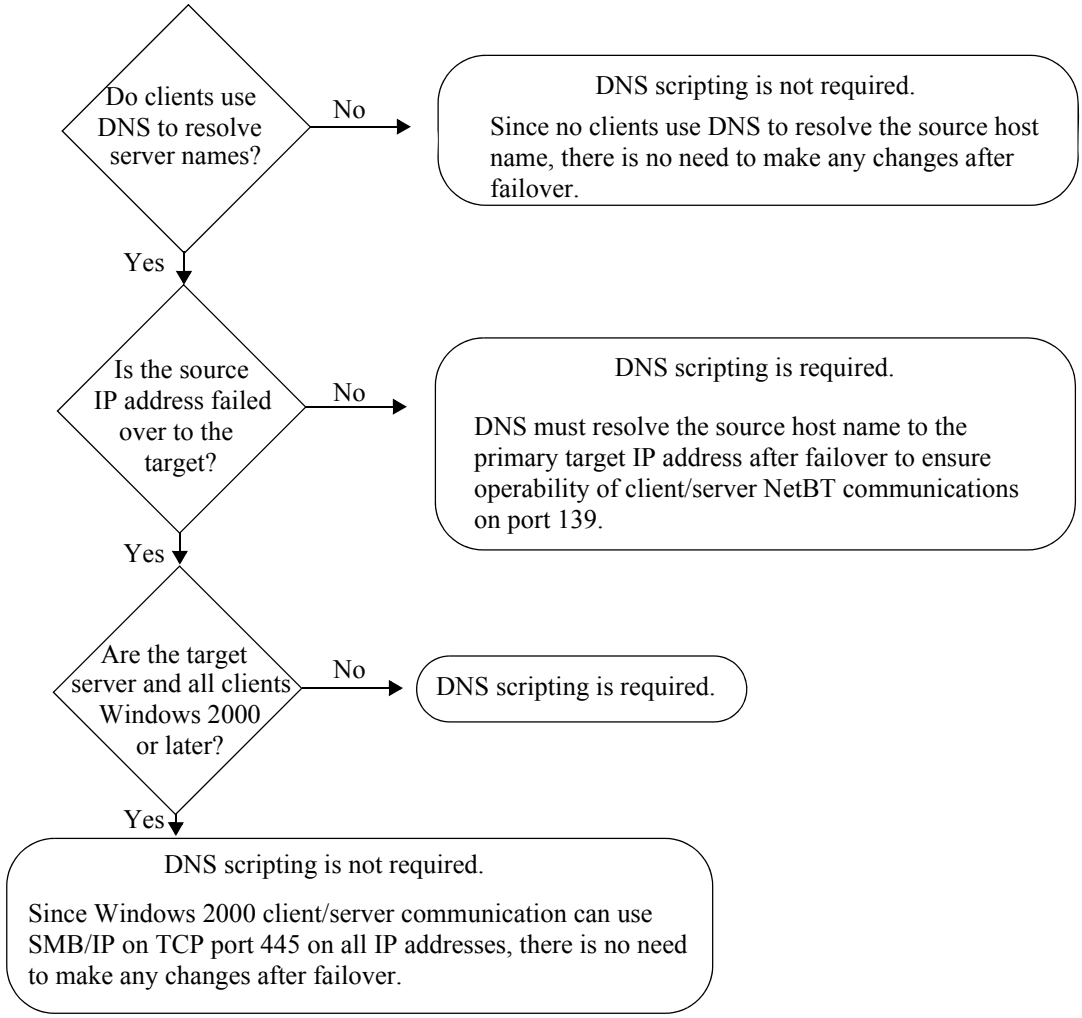

The DNS Server Troubleshooting Tool utility (DNSCMD) from the Windows 2000 Support Tools can be used in the failover and failback scripts to delete and add host and reverse lookup entries so that the source host name will resolve to the target IP address. The following example commands delete the host and reverse lookup entries associating the host name "europa" with 192.168.15.14 in the `nsisw.com` zone on the DNS server (`dnssvr.nsisw.com`), and add the entries to associate europa with 192.168.15.18.

```
dnscmd dnssvr.nsisw.com /RecordDelete nsisw.com europa A 192.168.15.14 /f
dnscmd dnssvr.nsisw.com /RecordDelete 168.192.in-addr.arpa 14.15 PTR
europa.nsisw.com /f
dnscmd dnssvr.nsisw.com /RecordAdd nsisw.com europa A 192.168.15.18
dnscmd dnssvr.nsisw.com /RecordAdd 168.192.in-addr.arpa 18.15 PTR
europa.nsisw.com
```

DNSCMD commands will only work if dynamic updates are enabled on the DNS zone. This is configured on the DNS zone's Properties dialog box in the Windows 2000 Microsoft Management Console DNS snap-in. If only secure updates is enabled (this option is available only on Active Directory-integrated zones), the DNSCMD utility must be used in the context of a user who is in the domain DnsAdmins group (i.e., the Storage Mirroring service logon account must be in the DnsAdmins group if the commands are in failover/failback scripts). In Windows 2003 Active Directory environments, the Storage Mirroring service logon account must be of a domain admin class in order to script DNSCMD. Failover/failback scripts do not run in the security context of the user specified in the failover monitor Account option despite the dialog box's implication (in Storage Mirroring 4.1) to the contrary.

The Windows 2000 Dynamic DNS (DDNS) client does not initiate a registration reflecting the failed-over name and IP address when failover occurs, and the `ipconfig /registerdns` command will not cause the failed-over name and IP address to be registered. Accordingly, host records for the source will remain intact after failover, and any required changes must be made on all DNS servers used by relevant clients.

Changes to non-Windows 2000 DNS servers and Windows 2000 DNS servers with dynamic updates disabled must be implemented by some other means. At this time, HP does not have any specific documentation on how to script changes to DNS records other than the Windows 2000 DDNS solution. However, since DNS zone files are text-based, they can be manipulated with any scripting language that can open, parse, and write to a text file.

Active Directory and Service Principal Names

Service Principle Names (SPNs) are properties of Active Directory (AD) computer accounts. When an AD client attempts to use a network share, it checks AD to see what computer is associated with the requested host name. When the client attempts to establish a session with the server, a message stating "Login Failure: The target account name is incorrect" will be returned if the computer account with which the SPN is associated does not belong to the server that receives the connection request.

When Storage Mirroring failover occurs, the source's SPNs must be deleted so that the target server will accept requests when clients attempt to access `\\SOURCE\SHARE`. If there are no SPNs associated with the name used in the request, the target server will allow the client connection since there is no conflict.

The "Write servicePrincipalName" permission on the source's AD computer account must be assigned to the account that will modify the SPNs. This is an advanced permission and assigning either of the more general Write or Full Control permissions, which are assigned to Domain Admins by default, will also be adequate. The permission must be assigned to one of the following:

- The target's Storage Mirroring service logon account. If the target's Storage Mirroring service is configured to log on as the System account, the target's Active Directory computer account should be assigned the permissions.
- The account specified in the failover monitor configuration.

Use the following steps to give an account the appropriate permissions to modify the source's SPNs.

1. Start Active Directory Users and Computers.
2. Select **View**, **Advanced**.
3. Locate the source's computer account.
4. Right-click on the source computer account and select **Properties**.
5. Select the Security tab and click **Advanced**.
6. If the account or group you want to add is not listed, click **Add** to add it.
7. Select the account or group and click **View/Edit**.
8. Select the Properties tab and check **Write servicePrincipalName**.
9. Click **OK** to accept the change.

There are two utilities that can be used to verify the SPN modifications and make the changes via command line if necessary. The NSISPN.EXE utility usage help is available by running NSISPN with no parameters. Microsoft's SETSPN.EXE utility is also available in the Windows 2000 Resource Kit, and has similar functionality and usage.

If the computer accounts have been moved from the default containers (the Computers and Domain Controllers containers), Storage Mirroring failover in versions prior to 4.2 Service Pack 1 and versions of the NSISPN utility prior to 1.1 may not make the necessary changes. Accordingly, please ensure that the version of NSISPN in use is 1.1 or later. The version is reported when NSISPN is run with no parameters.

After failover, the source SPNs can be viewed by running the following command:

```
NSISPN -L SOURCE_NAME
```

If the following SPNs are present, they must be deleted in order for clients to use the source name to access shares on the target:

```
HOST/SOURCE_NAME
HOST/SOURCE_NAME.domain.com
```

Following is a sample script to remove the source SPNs and add them to the target. This can be run as a Storage Mirroring failover script and may be necessary if the source computer is not in the default AD container when using versions of Storage Mirroring prior to 4.2 Service Pack 1. If NSISPN is used in a failover script, the Storage Mirroring service logon account must have the appropriate permissions to delete and add the SPNs since the failover scripts run in its security context. All instances of SOURCE and TARGET must be replaced with the associated computers' names, and domain.com should be replaced with the appropriate AD domain name. This script removes/adds both the HOST and SMTPSVC SPNs, which are typical for a Windows 2000 server running Exchange Server 5.5.

SAMPLE_FAILOVER.BAT

```
NSISPN -D HOST/SOURCE.domain.com SOURCE
NSISPN -D HOST/SOURCE SOURCE
NSISPN -D SMTPSVC/SOURCE.domain.com SOURCE
NSISPN -D SMTPSVC/SOURCE SOURCE

NSISPN -A HOST/SOURCE.domain.com TARGET
NSISPN -A HOST/SOURCE TARGET
NSISPN -A SMTPSVC/SOURCE.domain.com TARGET
NSISPN -A SMTPSVC/SOURCE TARGET
```

The following sample failback script will return the SPNs back to the original configuration.

SAMPLE_FAILBACK.BAT

```
NSISPN -D HOST/SOURCE.domain.com TARGET
NSISPN -D HOST/SOURCE TARGET
NSISPN -D SMTPSVC/SOURCE.domain.com TARGET
NSISPN -D SMTPSVC/SOURCE TARGET

NSISPN -A HOST/SOURCE.domain.com SOURCE
NSISPN -A HOST/SOURCE SOURCE
NSISPN -A SMTPSVC/SOURCE.domain.com SOURCE
NSISPN -A SMTPSVC/SOURCE SOURCE
```

If the target server is a domain controller, any SPNs added to its computer account will be removed periodically. This is a normal domain controller function and will not affect client access. However, since there are no SPNs for the source name, it may be possible for clients to be inadvertently or maliciously redirected and connected to another server while in a failover state.

Name caching

By default, Windows systems cache DNS and NetBIOS name resolutions in the DNS Resolver Cache and NetBIOS Remote Cache Name Table respectively. This functionality does not impede the ability of clients to access the source name after failover as long as the appropriate WINS and DNS changes are made. Even in name-only failover scenarios (where the IP address is not failed over), clients will use WINS or DNS to re-resolve the name if an attempt to initialize a session with the cached entry fails.

Although name caching does not present any issues if WINS and DNS are updated after failover, information about the entries in the name cache may be useful in troubleshooting if difficulties are experienced. The cached entries can quickly show whether the client has resolved the source name to the correct IP address.

A system's NetBIOS Remote Cache Name Table can be viewed by using the `nbtstat -c` command, and the `ipconfig /displaydns` command displays the DNS Resolver Cache. The NetBIOS and DNS name caches can be purged with the nbtstat -R (the -R is case-sensitive) and `ipconfig /flushdns` commands respectively.

Microsoft Knowledge Base articles 120642, 245437, and 187709 contain information regarding the configuration of timeout values for the NetBIOS and DNS name caches. However, keep in mind that the name resolution cache settings do not need to be adjusted for successful failover. If WINS and DNS entries

are updated properly, previously cached name resolutions will not impede the ability of clients to establish SMB sessions with the target server.

Storage Mirroring failover and ARP

After TCP/IP has resolved the host name to an IP address, it passes the IP packet to ARP (Address Resolution Protocol). ARP resolves the IP address to an adapter MAC (media access control) address and then passes the packet to the data-link layer (Ethernet, Token Ring, ATM, etc.). ARP maintains a cache of IP address-to-MAC address resolutions on each system. This cache must be updated at failover so that clients with cached entries will not attempt to send packets to the source's MAC address.

When Storage Mirroring is installed, the ARP Responder device driver is installed and set to a startup type of demand (Windows 2000) or manual (Windows NT 4.0). Storage Mirroring uses the ARP Responder to broadcast an unsolicited (gratuitous) ARP when failover occurs. The unsolicited ARP forces the systems on the same physical network to update their ARP caches with an entry associating the source IP address with the target adapter's MAC address. The following event will be created in the target's Application Event Log when the unsolicited ARP is broadcast:

```
Event ID: 5400
Source: Storage Mirroring
Type: Information
Description: Broadcasted new MAC address [target adapter MAC address] for IP
address [source IP address].
```

The `arp -a` command can be used to see the ARP cache on a system for troubleshooting purposes.

Confirming Storage Mirroring failover

When Storage Mirroring executes a failover, it makes entries in the Storage Mirroring logs, the Windows Event Log, and will generate an SNMP trap (if the SNMP component is installed). Operating system commands can also be used to show that the appropriate name and IP address changes have been made.

The following messages are logged in the target's Storage Mirroring log file when failover occurs:

```
05/01/200213:34:14.5230101Failover in progress!!!
05/01/200213:34:44.2860700001Failover complete for SERVER1
```

A number of entries will be made in the Application Event Log as well:

```
Event ID: 5100
Source: Storage Mirroring
Type: Information
Description: Failover completed for [source computername]
Event ID: 5101
Source: Storage Mirroring
Type: Information
Description: IP address [source IP address] with subnet mask [source subnet mask]
was added to target machine's [target adapter name] adapter.
```

If the SNMP service and the Storage Mirroring SNMP component are installed, the target will generate the following SNMP trap when failover occurs:

```
Trap: DttrapFailoverInProgress
Description: Failover is occurring.
```

The `nbtstat -n` command can be used to verify that the source's name is successfully added to the target after failover. This command shows the names registered on the local system. Following is a sample that shows `nbtstat -n` output on the target server CALLISTO before and after failover of the source server GANYMEDE (in the JUPITER domain).

Before Failover

```
D:\>nbtstat -n
TS:
Node IpAddress: [172.16.137.31] Scope Id: []

                NetBIOS Local Name Table
       Name               Type         Status
    ---------------------------------------------
    CALLISTO       <00>  UNIQUE      Registered
    JUPITER        <00>  GROUP       Registered
    CALLISTO       <20>  UNIQUE      Registered
    JUPITER        <1E>  GROUP       Registered
    INet~Services  <1C>  GROUP       Registered
    IS~callisto....<00>  UNIQUE      Registered
```

After Failover

```
D:\>nbtstat -n
TS:
Node IpAddress: [172.16.137.31] Scope Id: []

                NetBIOS Local Name Table
       Name               Type         Status
    ---------------------------------------------
    CALLISTO       <00>  UNIQUE      Registered
    JUPITER        <00>  GROUP       Registered
    CALLISTO       <20>  UNIQUE      Registered
    JUPITER        <1E>  GROUP       Registered
    INet~Services  <1C>  GROUP       Registered
    IS~callisto....<00>  UNIQUE      Registered
    GANYMEDE       <00>  UNIQUE      Registered
    GANYMEDE       <03>  UNIQUE      Registered
    GANYMEDE       <20>  UNIQUE      Registered
```

Additionally, `nbtstat -a [target IP address]` can be run from a remote system to show the target's name table to retrieve the same information.

The `ipconfig` command can be used to verify that the source's IP address has failed over to the target. Failed-over IP addresses will be listed in the ipconfig output as well as in the network adapter's TCP/IP Properties Advanced dialog box.

Replace option

Storage Mirroring includes a replace option that enables the target to replace its computer name and IP address with the source's computer name and IP address at failover. This option is supported on Windows NT 4.0 targets in the released versions of Storage Mirroring, and in a hotfix version for Windows 2000.

When the replace option is used on a Windows NT 4.0 target, the Computer Browser, Net Logon, and Server services are stopped and started after the name and IP address is changed when failover occurs. Windows 2000 systems must be rebooted after the failover occurs if the replace option is used.

Since the source name and IP address remain the same when the replace option is used, there are no name resolution issues to consider as long as the source's IP address is the target's primary IP address after failover occurs.

Storage Mirroring failover and domain controllers

HP recommends that solutions using Storage Mirroring failover be implemented on member servers whenever possible. However, there are certain environments where the use of a domain controller as the source or target is unavoidable. Domain controllers can be successfully failed over with Storage Mirroring when necessary, but the domain controller functionality is not failed over. Following is a discussion of the

issues that should be considered when implementing a Storage Mirroring failover solution with domain controllers.

Windows 2000

Windows 2000 Active Directory domain controllers use a pull-based replication architecture, so there is no risk of Active Directory updates being sent to the wrong server due to Storage Mirroring failover adding the source's computer name to the target. The only items to consider are the effects of a given domain controller being unavailable. A brief outline of some of these issues follows, but a complete understanding can only be gained by an in-depth knowledge of Active Directory. Active Directory documentation is included in the Windows 2000 Resource Kit.

Active Directory has five FSMO (Flexible Single Master Operation) roles, and each role is assigned to one domain controller in the domain or forest. Role ownerships can be easily moved between domain controllers to facilitate changes to the domain controller infrastructure. The five FSMO roles are:

1. Schema master (one per forest)
2. Domain naming master (one per forest)
3. Primary domain controller (PDC) emulator (one per domain)
4. Routing information daemon (RID) master (one per domain)
5. Infrastructure master (one per domain)

See Microsoft Knowledge Base article 197132 for a concise description of these roles. Unavailability of some of the FSMO roles can cause immediate effects, such as Windows NT 4.0 users not being able to change their passwords (PDC emulator), inability to extend the AD schema (schema master), and inability to add a domain to a forest (domain naming master).

Global Catalog (GC) servers are also critical for proper domain functionality, particularly the logon process in a multi-domain forest (see Microsoft Knowledge Base article 216970). A properly designed Active Directory infrastructure will have multiple GC servers placed strategically throughout the network to ensure that the failure of a given GC server will not impact users.

Windows NT

Windows NT domains also use a pull-based directory replication architecture, so a failed-over backup domain controller (BDC) will not cause any inconsistencies in directory replication. The PDC prompts BDCs to request replication on a scheduled interval, and the BDCs then request updates from the PDC. The BDC's request informs the PDC of the last change it received, and the PDC sends the subsequent updates.

Again, there are no issues related to Storage Mirroring, and the only concerns are those related to a PDC or BDC being unavailable for a period of time. A thorough understanding of Windows NT domains will be necessary to be aware of all possible issues, and the Windows NT 4.0 Resource Kit contains the relevant documentation. Some of the most important issues to consider include:

- Clients may be required to authenticate across a WAN link if a local domain controller is not available, which may cause a delay based on the available bandwidth.
- Extended downtime of the BDC may result in a full directory synchronization, which, depending on the size of the directory, can utilize significant bandwidth. By default, the PDC keeps a change log of 2000 entries, and a BDC will require a full synchronization if more than 2000 changes are made during its downtime. This is particularly a concern if the BDC is separated from the PDC by a WAN link.
- Some applications, such as Microsoft Exchange Server 5.5, require access to the PDC at installation time.
- User and computer accounts, user rights, and other directory objects cannot be created, modified, or deleted if the PDC is unavailable.
- Trusts cannot be created if the PDC is unavailable.

IP address failover to a remote target

In most cases, failing over to a remote target can be accomplished by failing over the computer name only and updating the name resolution servers to associate the source name with the target's IP address. However, some solutions may require an IP address to be failed over to a remote target. Failing over IP addresses to remote targets can be accomplished a number of ways. If a VPN infrastructure exists so that

the source and target can be on the same subnet, IP address failover will work exactly as it does in a LAN environment.

If a VPN does not exist, routers can be automatically reconfigured with a failover script after the IP address is failed over so that the failed-over address will be routed to the target. This would entail configuring the routers to move the source's subnet from the source's physical network to the target's physical network. There are a number of issues to consider when designing a solution that requires router configuration to achieve IP address failover. Since the route to the source's subnet will be changed at failover, the source server must be the only system on that subnet, which in turn requires all server communications to pass through a router. Additionally, it may take several minutes or even hours for routing tables on other routers throughout the network to converge.

Depending on the router's capabilities, other options may also exist. Some routers can be configured to provide a routing infrastructure that can accommodate IP address failover to another segment. Additional discussion on this topic is beyond the scope of this document due to the number of router manufacturers and various capabilities of router operating systems.

Troubleshooting client access after failover

In general, the recommended way to resolve client access issues is to confirm that everything is configured appropriately. This can be done by gathering information about the solution requirements and the network infrastructure. The following information, at a minimum, is required in order to design the failover configuration:

- Which clients will be accessing the failed-over server?
- What are the client, source, and target operating systems?
- Where are the clients and target located?
- What are the name resolution method of clients?
- What is the WINS configuration of clients and target?
- What is the DNS configuration of the clients?
- What is the WINS architecture?
- What is the DNS architecture?
- What type of domain is in use (Windows NT or Active Directory)?
- Where are the Active Directory domain controllers located?
- In what container is the source's Active Directory computer account?

A failover solution may work in a test environment but fail in production if not properly designed, and the information listed above is required in order to design it properly. Until the solution is properly designed, actions taken to resolve a specific issue may appear to have resolved an issue when in fact they were irrelevant. For example, if a Windows NT 4.0 client is using a WINS server that did not get updated at failover, it will not be able to access the failed-over server until its WINS server receives replication from the target's WINS server. If the Storage Mirroring service account is changed and failover is tried again, the client may be able to access the failed-over server, not because of the change that was made, but because the WINS replication finally occurred just before failback. Accordingly, a “shotgun” approach to resolving issues may lead to false conclusions, confusion, and extended troubleshooting time.

In general, failure to update WINS servers and SPNs are the most common causes of client access issues after failover. If the client or target is not Windows 2000 or later, then the issue is most likely failure to update WINS servers. If the client and target are Windows 2000 or later, then the issue is most likely failure to update SPNs.

After the correct configuration has been determined and verified, troubleshooting should begin with verifying that the IP address and name were failed over correctly. This can be done by running the following two commands:

```
Nbtstat -n
IPConfig /all
```

Caches containing MAC address, NetBIOS, and DNS host names have not been known to cause any issues with client access and flushing the caches will not resolve the issue. In general, if the requested

operation fails, the operation is retried without using the cached entry. Viewing the cached entries is usually only relevant to finding out what the name resolution servers are returning.

Diagnosing common issues

Issue	Clients receive “The network path was not found” when attempting to access the failed-over source.
Diagnosis	This error is returned if the source name is resolved to an IP address or TCP/IP port that is not responding.
Possible Causes	1. The client or target is pre-Windows 2000 and the name resolution server (WINS or DNS) used by the client is not being updated to associate the source name with the target's primary IP address. Failure to update the WINS server is by far the most common cause of this issue. 2. The source's IP address or name was not added to the target at failover. 3. The client is pre-Windows 2000 and NetBIOS over TCP/IP (NetBT) is disabled on the target network adapter.
Primary Troubleshooting	1. Confirm that the failover monitor is configured to fail over the source's name. 2. Confirm that the failover monitor is configured with failover scripts to update all WINS and DNS servers used by all clients to associate the source name with the target's primary IP address. 3. Confirm that the Storage Mirroring service account has the necessary permissions to complete the failover script commands that make the changes to the WINS and DNS servers. 4. If clients are pre-Windows 2000, confirm that NetBT is enabled on the appropriate target network adapter.
Advanced Troubleshooting	1. Run `nbtstat -n` on the target to confirm that the source's name was failed over. If the source name does not appear in the NetBIOS local name table, check the failover monitor configuration to ensure that it is configured to fail over the source name. Check the Windows application log and Storage Mirroring log on the target to confirm that failover occurred. 2. Check the WINS records associated with the source name on all WINS servers used by clients. Check the following items if the WINS entries do not associate the source name to the target's primary IP address. a. Confirm that failover scripts have been configured to update the WINS servers. b. Confirm that the Storage Mirroring service account has the appropriate permissions to successfully complete the commands to update the WINS servers. This can be done by logging on to the target server with the Storage Mirroring service account, running WINSCL or NETSH interactively, and entering the commands manually. This can also be done by redirecting the output of the failover script to a text file to review the script results. c. Confirm that the failover scripts are correct by running WINSCL or NETSH interactively and entering the commands manually. This can also be done by redirecting the output of the failover script to a text file to review the script results. 3. Run `IPCONFIG /ALL` on the target to confirm that the source's IP address was added to the correct adapter.

Issue	Windows 2000 and later clients receive “login failure, the target account name is incorrect” or “access denied” when attempting to access a failed-over source.
Diagnosis	Service Principal Names are properties of Active Directory computer accounts. When clients make a connection to a server, they query Active Directory to find the HOST SPNs for the server name used in the connection request. If a HOST SPN exists, the connection will fail if the server receiving the connection request is not the computer associated with the computer account that owns the SPN.

Possible Causes	1. SPNs are not getting changed at failover. 2. Clients are using a domain controller that has not yet received the changes to the SPNs.
Primary Troubleshooting	1. Confirm that the failover monitor is configured to fail over and fail back the Active Directory host name. 2. If NSISPN is being used in failover scripts, confirm that it is version 1.1 or later.
Advanced Troubleshooting	1. After failover, run "`NSISPN -L [source]`" from the target and save the output for review. 2. If the `HOST/SOURCE` and/or `HOST/SOURCE.domain.com` SPNs are present, confirm the following: a. The "failover hostname" and "failback hostname" options are enabled on the failover monitor. b. If Storage Mirroring 4.2 Service Pack 1 or later is not in use and the source's Active Directory computer account is not in the Computers container, use the NSISPN (version 1.1 or later) utility in the failover scripts to make the necessary SPN changes. c. Ensure that the "Write servicePrincipalName" permission (at a minimum) has been assigned to the appropriate account, as explained in "Active Directory and Service Principal Names" on page 10. 3. If the HOST SPNs are not present, run "`NSISPN -L [source]`" from the client. If the SPNs are different from those reported when running the same command from the target, then the issue is most likely due to the lag in propagation of the SPN changes to other domain controllers. This is most likely to occur when there are multiple sites in AD and the client is in a different site. This can be resolved by forcing replication after failover with the Replication Diagnostics Tool (`Repadmin.exe`) in the Windows 2000 Support Tools. The Replication Diagnostics Tool can be used in the failover script to ensure that changes to the SPNs are replicated to other domain controllers immediately.

Issue	Clients cannot print after failover occurs.
Diagnosis	The Windows NT 4.0 and Windows 2000 print spooler binds to the name that is specified as the ActiveComputerName, and requests to the spooler that use a different name will fail. The Storage Mirroring ChngName utility must be used in the failover script to change the ActiveComputerName. See the *Guidelines for Using Storage Mirroring with Windows Print Servers* application note, available from the HP web site, for details on using the ChngName utility for print server failover
Possible Causes	1. The ChngName utility was not used correctly. 2. A name resolution or SPN issue is causing attempted client connections to fail. 3. Failover for more than one source has occurred, and printing is only working for the last source that was failed over. If failover has occurred for more than one source, printing will only be available to the clients that were using the last source to failover. The print spooler will only respond to requests that use the ActiveComputerName that was in place when the print spooler service was last started.
Primary Troubleshooting	1. Confirm that the ChngName utility was used correctly in the failover scripts. `ChngName /t` must not be run in the failover script for Windows 2000 print servers. 2. Try to view a share on the target using the source's name. If clients are getting "Network path not found," or other error messages, see the previous sections to resolve name resolution or SPN issues. 3. Run the failover script manually to ensure that all commands complete correctly and resolve any failures.

Issue	NFS clients cannot access mounts after failover.
Diagnosis	UNIX clients may see NFS mounts become unresponsive after failover. This is a limitation of the NFS client software and cannot be addressed with any Storage Mirroring configuration. Some clients will handle the failover event seamlessly by automatically reconnecting the session. See the appropriate documentation or contact the NFS client software publisher's technical support to determine whether the client can be configured to automatically reconnect.